COOPERATIVE LEARNING IN MULTI-AGENT SYSTEMS FROM INTERMITTENT MEASUREMENTS*

NAOMI EHRICH LEONARD† AND ALEX OLSHEVSKY‡

Abstract. Motivated by the problem of decentralized direction-tracking, we consider the general problem of cooperative learning in multi-agent systems with time-varying connectivity and intermittent measurements. We propose a distributed learning protocol capable of learning an unknown vector μ from noisy measurements made independently by autonomous nodes. Our protocol is completely distributed and able to cope with the time-varying, unpredictable nature of inter-agent connectivity, repeated failures of nodes, and intermittent noisy measurements of μ. Our main results bound the learning speed of our protocol in terms of a novel measure of graph connectivity we call the sieve constant of a graph.

Key words. multi-agent systems, learning theory, distributed control.

AMS subject classifications. 93E35, 93A14.

1. Introduction.

Widespread deployment of mobile sensors is expected to revolutionize our ability to monitor and control physical environments. However, for these networks to reach their full range of applicability they must be capable of operating in uncertain and unstructured environments. Realizing the full potential of networked sensor systems will require the development of protocols that are fully distributed and adaptive in the face of persistent faults and time-varying, unpredictable environments.

Our goal in this paper is to initiate the study of cooperative multi-agent learning by distributed networks operating in unknown and changing environments, subject to node faults and failures of communication links. While our focus here is on the basic problem of learning an unknown vector, we hope to contribute to the development of a broad theory of cooperative, distributed learning in such environments, with the ultimate aim of designing sensor network protocols capable of learning and adaptability.

We will study a simple, local protocol for learning a vector from intermittent measurements and evaluate its performance in terms of the number of nodes and the (time-varying) network structure. Our direct motivation is the problem of direction tracking from chemical gradients. A network of mobile sensors needs to move in a direction μ (understood as a vector on the unit circle), which none of the sensors initially knows; however, intermittently some sensors are able to obtain a noisy sample of μ. The sensors can observe the velocity of neighboring sensors but, as the sensors move, the set of neighbors of each sensor changes; moreover, new sensors occasionally join the network and current sensors sometimes permanently leave the network. The challenge is to design a protocol by means of which the sensors can adapt their velocities based on the measurements of μ and observations of the velocities of neighboring sensors so that every node's velocity converges to μ as fast as possible.

We will consider a natural generalization in the problem, wherein we abandon the constraint that μ lies on the unit circle and instead consider the problem of learning

*This work was supported in part by AFOSR grant FA9550-07-1-0-0528, ONR grant N00014-09-1-1074 and ARO grant W911NG-11-1-0385.

†Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ, 08544, USA (naomi@princeton.edu).

‡Department of Industrial and Enterprise Systems Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, 61801, USA (aolshev2@illinois.edu).

an arbitrary vector μ by a network of mobile nodes subject to time-varying (and unpredictable) inter-agent connectivity, repeated failures of nodes, and intermittent, noisy measurements. We will be interested in the speed at which local, distributed protocols are able to drive every node's estimate of μ to the correct value. We will be especially concerned with identifying the salient features of network topology that result in good (or poor) performance.

1.1. Cooperative multi-agent learning. We begin by formally stating the problem for a fixed number of nodes. We consider n autonomous nodes engaged in the task of learning a vector $\mu \in \mathbb{R}^k$. At each time $t = 0, 1, 2, \ldots$ we denote by $G(t) = (V(t), E(t))$ the graph of inter-agent communications at time t: two nodes are connected by an edge in $G(t)$ if and only if they are able to exchange messages at time t. Note that by definition the graph $G(t)$ is undirected. If $(i, j) \in G(t)$ then we will say that i and j are neighbors at time t. We will adopt the convention that $G(t)$ contains no self-loops. While we do not assume that the graphs $G(t)$ are connected at each time t, we do assume they satisfy a standard condition of uniform connectivity: there exists some constant positive integer B (unknown to any of the nodes) such that the graph sequence $G(t)$ is B-connected, i.e. the graphs $(\{1, \ldots, n\}, \cup_{kB+1}^{(k+1)B} E(t))$ are connected for each integer $k \geq 0$. Intuitively, the uniform connectivity condition means that once we take all the edges that have appeared between times kB and $(k+1)B$, the graph is connected.

Each node maintains an estimate of μ; we will denote the estimate of node i at time t by $v_i(t)$. At time t, node i can update $v_i(t)$ as a function of the vaues $v_j(t)$ held by all neighbors j of node i. Physically, these updates may be the result of a message exchange or may come about through observations by each node. Occasionally, some nodes have access to a noisy measurement

$$\mu_i(t) = \mu + w_i(t),$$

where $w_i(t)$ is a zero-mean random vector every entry of which has variance σ^2; we assume all vectors $w_i(t)$ are independent of each other and of all $v_i(t)$. In this case, node i incorporates this measurement into its updated estimate $v_i(t+1)$. We will make an assumption of uniform measurement speed, namely that at least one node has access to a measurement every T steps.

It is useful to think of this formalization in terms of our motivating scenario, which is a collection of nodes - vehicles, UAVs, mobile sensors, or underwater gliders - which need to learn and follow a direction. Updated information about the direction arrives from time to time as one or more of the nodes takes measurements, and the nodes need a protocol by which they update their velocities $v_i(t)$ based on the measurements and observations of the velocities of neighboring nodes.

This formalization also describes the scenario in which a moving group of animals must all learn which way to go based on intermittent samples of a preferred direction and social interactions with near neighbors. An example is collective migration where high costs associated with obtaining measurements of the migration route suggest that the majority of individuals rely on the more accessible observations of the relative motion of their near neighbors when they update their own velocities $v_i(t)$ [15].

1.2. Our results. We now describe the learning protocol which we analyze for the remainder of this paper. If at time t node i does not have a measurement of μ, it

moves its velocity in the direction of its neighbors:

$$v_i(t+1) = v_i(t) + \frac{\Delta(t)}{4} \sum_{j \in N_i(t)} \frac{v_j(t) - v_i(t)}{\max(d_i(t), d_j(t))}. \tag{1.1}$$

Here $N_i(t)$ is the set of neighbors of node i, $d_i(t)$ is the cardinality of $N_i(t)$, and $\Delta(t)$ is a stepsize which we will specify later.

On the other hand, if node i does have a measurement $\mu_i(t)$, it updates as

$$v_i(t+1) = v_i(t) + \frac{\Delta(t)}{4} \left(\mu_i(t) - v_i(t)\right) + \frac{\Delta(t)}{4} \sum_{j \in N_i(t)} \frac{v_j(t) - v_i(t)}{\max(d_i(t), d_j(t))}. \tag{1.2}$$

Intuitively, each node seeks to align its estimate $v_i(t)$ with both the measurements it takes and the estimates of neighboring nodes. As nodes align with one another, information from each measurement slowly propagates throughout the system.

Our protocol is motivated by a number of recent advances within the literature on multi-agent consensus. On the one hand, the weights we accord to neighboring nodes are based on Metropolis weights (first introduced within the context of multi-agent control in [6]) and are chosen because they lead to a tractable Lyapunov analysis as in [22]. On the other hand, we introdue a stepsize $\Delta(t)$ which we will later choose to decay to zero with t at an appropriate speed by analogy with the recent work on multi-agent optimization [23, 28, 32].

Our protocol is also motivated by models used to analyze collective decision making and collective motion in animal groups [13, 17]. Our time varying stepsize rule is similar to models of context-dependent interaction in which individuals reduce their reliance on social cues when they are progressing towards their target [30].

We now state our main results. We will show that this distributed learning protocol converges to the correct answer under fairly minimal assumptions and derive upper bounds on the convergence rate. The first theorem states the basic convergence result which underpins the subsequent results.

THEOREM 1.1. *If the stepsize $\Delta(t)$ is nonnegative, nonincreasing and satisfies*

$$\sum_{t=1}^{\infty} \Delta(t) = +\infty, \qquad \sum_{t=1}^{\infty} \Delta^2(t) < \infty, \qquad \sup_{t \geq 1} \frac{\Delta(t)}{\Delta(t+c)} < \infty \qquad \textit{for any integer } c$$

then for any initial values $v_1(0), \ldots, v_n(0)$, we have that with probability 1

$$\lim_{t \to \infty} v_i(t) = \mu \qquad \textit{for all } i.$$

Our main goal in this paper is to prove strengthened versions of Theorem 1.1 which provide quantitative bounds on the rate at which convergence to μ takes place. We are particularly interested in the scaling of the convergence time with the combinatorics of structure of the interconnection graphs. We will adopt the natural measure of how far we are from convergence, namely the sum of the squared distances from the final limit:

$$Z(v(t)) = \sum_{i=1}^{n} ||v_i(t) - \mu||_2^2.$$

The next theorem provides an upper bound on convergence time for a particularly chosen stepsize in the case when $B = 1$, i.e., when each graph $G(t)$ is connected.

THEOREM 1.2. *If $B = 1$ and $\Delta(t) = 1/t^{1-\epsilon}$ for any $\epsilon \in (0,1)$ then*

$$\limsup_{t\to\infty} t^{1-\epsilon} E[Z(v(t)) \mid v(0)] \leq \frac{n\sigma^2 T^{1-\epsilon} Z(0)}{2\min_{t=1,2,\ldots} \kappa(G(t))}, \tag{1.3}$$

where $\kappa(G)$ is a measure of graph connectivity which we call the sieve constant *of a graph, defined as follows. For a nonnegative, stochastic matrix $A \in \mathbb{R}^{n\times n}$, the sieve constant $\kappa(A)$ is defined as*

$$\kappa(A) = \min_{m=1,\ldots,n} \min_{||x||_2=1} \; x_m^2 + \sum_{k\neq l} a_{kl}(x_k - x_l)^2.$$

For an undirected graph $G = (V, E)$, the sieve constant $\kappa(G)$ denotes the sieve constant of the Metropolis matrix, which is the stochastic matrix with

$$a_{ij} = \begin{cases} \frac{1}{\max(d_i,d_j)}, & \text{if } (i,j) \in E \text{ and } i \neq j, \\ 0, & \text{if } (i,j) \notin E. \end{cases}$$

To parse the statement of Theorem 1.2, observe that the right-hand side of Eq. (1.3) does not depend on t. Consequently, assuming a positive but very small $\epsilon \approx 0$ is chosen, the theorem states that the expected squared error $E[Z(v(t)) \mid v(0)]$ asymptotically decays nearly as fast $1/t$. The bulk of the theorem statement describes the constant in front of the nearly linear decay, which depends on the problem parameters, namely the number of nodes n, the sampling period T, the variance of the noise σ^2, and the graph interconnection sequence $G(t)$. Crucially, the only influence that the structure of the graphs $G(t)$ has on this bound is in terms of the inverse of sieve constants $\kappa(G(t))$.

This naturally raises the question of how large $1/\kappa(G)$ can be for an undirected graph G on n nodes, and how it depends on the connectivity properties on the graph G. We answer this question in the following theorem.

THEOREM 1.3. *For any undirected connected graph G,*

$$\frac{1}{\kappa(G)} \leq n d_{\max} D,$$

where $d_{\max}$ is the largest degree of a node in G and D is the diameter of G.

Taken together, the previous two theorems imply upper bounds on the amount of time it takes for $E[Z(v(t)) \mid v(0)]$ to shrink below $\epsilon Z(0)$. Crucially, these upper bounds are polynomial in terms of the number of nodes n. We note, also, that our bounds can be much better (potentially by orders of magnitude) when diameter is low and degrees are small. Intuitively, a low diameter ensures that information propagates through few hops from node to node while lower degrees ensure that new pieces of information have more influence in nearest-neighbor interactions.

The preceeding theorem admits a generalization to the case of general B, in which case the convergence rate is naturally bounded in terms of the sieve constant of the line graph.

THEOREM 1.4. *If $\Delta(t) = 1/t^{1-\epsilon}$ then*

$$\limsup_{t\to\infty} t^{1-\epsilon} E[Z(v(t)) \mid v(0)] \leq \frac{4nd_{\max}\sigma^2 \max(T,B)^{2-\epsilon} Z(0)}{\kappa(L_n)},$$

where L_n is the undirected line graph on n nodes.

We will later sketch how these theorems may be adapted to the setting in which the number of nodes is not fixed but rather nodes can join and permanently depart from the system.

1.3. Related work. We believe that our paper is the first to derive rigorous results for the problem of cooperative multi-agent learning by a network subject to node failures, communication disruptions, and intermittent measurements. The key features of our model are 1) its cooperative nature (many nodes working together) 2) its reliance only on distributed and local observations 3) the incorporation of time-varying communication restrictions 4) its flexibility with respect to node faults (our results carry over to the setting in which nodes are allowed to enter and leave the system). These features are typically required in modern cyber-physical systems.

Naturally, our work is not the first attempt to fuse learning algorithms with distributed control or multi-agent settings. Indeed, the study of learning in games is a classic subject which has attracted considerable attention within the last couple of decades due in part to its applications to multi-agent systems. We refer the reader to the recent papers [2, 4, 5, 10, 9, 7, 14, 21, 1, 20] as well as the classic works [18, 12] which study multi-agent learning in a game-theoretic context. Moreover, the related problem of distributed reinforcement learning has attracted some recent attention; we refer the reader to [18, 29, 26]. We make no attempt to survey these literatures here and refer the reader to the references in the above papers, as well as the surveys [27, 24].

Finally, we note that much of the recent literature in distributed robotics has focused on distributed algorithms robust to faults and communication link failures. The number of works is once again too vast to survey, but we refer the reader to the representative papers [3, 19]. Our work here is very much in the spirit of that literature.

1.4. Outline. The remainder of this paper is organized as follows. Theorems 1-4 are proved in Section 2; moreover, Section 2 concludes with an extended remark sketching how our results can be carried over to settings where the number of nodes itself varies with time. We compute the sieve constant for a variety of graphs in Section 3. Section 4 provides some practical simulations of our learning protocol and Section 5 concludes with a summary of our results and a list of several open questions.

2. Proofs of the main results. The purpose of this section is to prove Theorems 1-4. In the course of this, we will demonstrate how the sieve constant naturally appears in the analysis of our learning protocol. We begin with some preliminary definitions.

2.1. Definitions. Given a nonnegative matrix $A \in \mathbb{R}^{n\times n}$, we will use $G(A)$ to denote the graph whose edges correspond to the positive entries of A in the following way: $G(A)$ is the directed graph on the vertices $\{1, 2, \ldots, n\}$ with edge set $\{(i,j) \mid a_{ji} > 0\}$. Note that if A is symmetric then the graph $G(A)$ will be undirected. We will use the standard convention of $\mathbf{e}_i$ to mean the i'th basis column vector and $\mathbf{1}$ to mean the all-ones vector. Finally, we will use $r_i(A)$ to denote the row

sum of the i'th row of A^2 and $R(A) = \mathrm{diag}(r_1(A), \ldots, r_n(A))$. When the argument matrix A is clear from context, we will simply write r_i and R for $r_i(A), R(A)$.

2.2. A few preliminary lemmas. In this subsection we prove a few lemmas which we will find useful in the proofs of our main theorems. Our first lemma gives a decomposition of a symmetric matrix and its immediate corollary provides a way to bound the change in norm arising from multiplication by a symmetric matrix. Similar statements were proved in [6],[22], and [31].

LEMMA 2.1. *For any symmetric matrix A,*

$$A^2 = R - \sum_{k<l} [A^2]_{kl} (\mathbf{e}_k - \mathbf{e}_l)(\mathbf{e}_k - \mathbf{e}_l)^T.$$

Proof. Observe that each term $(\mathbf{e}_k - \mathbf{e}_l)(\mathbf{e}_k - \mathbf{e}_l)^T$ in the sum on the right-hand side has row sums of zero, and consequently both sides of the above equation have identical row sums. Moreover, both sides of the above equation are symmetric. This implies it suffices to prove that all the (i,j)-entries of both sides with $i<j$ are the same. But on both sides, the (i,j)'th element when $i<j$ is $[A^2]_{ij}$. ☐

This lemma may be used to bound how much the norm of a vector changes after multiplication by a symmetric matrix.

COROLLARY 2.2. *For any symmetric matrix A,*

$$||Ax||_2^2 \;=\; ||x||_2^2 - \sum_{j=1}^{n} (1-r_j) x_j^2 + \sum_{k<l} [A^2]_{kl} (x_k - x_l)^2.$$

Proof. By Lemma 2.1,

$$\begin{aligned} ||Ax||_2^2 &= x^T A^2 x \\ &= x^T R x - \sum_{k<l} [A^2]_{kl} x^T (\mathbf{e}_k - \mathbf{e}_l)(\mathbf{e}_k - \mathbf{e}_k)^T x \\ &= \sum_{j=1}^{n} r_j x_j^2 - \sum_{k<l} [A^2]_{kl} (x_k - x_l)^2. \end{aligned}$$

Thus the decrease in squared norm from x to Ax is

$$||x||_2^2 - ||Ax||_2^2 = \sum_{j=1}^{n} (1-r_j) x_j^2 + \sum_{k<l} [A^2]_{kl} (x_k - x_l)^2.$$

☐

We now prove a basic positivity property for the sieve constant of a stochastic matrix (defined in Section 1), which we will have occasion to use in the next subsection.

LEMMA 2.3. *$\kappa(A) \geq 0$ and $\kappa(A) > 0$ if and only if the graph $G(A)$ is weakly connected*[1]*.*

Proof. It is evident from the definition of $\kappa(A)$ that it is necessarily nonnegative. If $G(A)$ is not weakly connected, then we can pick m as any vertex, set $x_i = 0$

[1] A directed graph is weakly connected if the undirected graph obtained by ignoring the orientations of the edges is connected.

on the connected component containing m, and $x_i = c$ on all the other connected components; here c is a positive constant chosen so that the normalization condition $||x||_2 = 1$ is satisfied. These choices of m and x result in $\kappa(A) = 0$.

On the other hand, suppose that $G(A)$ is weakly connected. Note that $\kappa(A) = 0$ implies that $x_m = 0$ and that every pair (i, j) with $a_{ij} > 0$ or $a_{ji} > 0$ satisfies $x_i = x_j$. Now the weak connectivity of $G(A)$ implies that every entry of x must be identical. Since $x_m = 0$, we in fact have that x is the zero vector, which contradicts the normalization condition $||x||_2 = 1$. Consequently, $\kappa(A) = 0$ is not possible if $G(A)$ is weakly connected. □

2.3. The sieve constant and the learning protocol. With the above lemmas in place, we can now turn to the analysis of our learning protocol. For the remainder of Subsection 2.3, we will assume that $k = 1$, i.e, μ and all $v_i(t)$ belong to $\mathbb{R}$. We will then define $v(t)$ to be the vector that stacks up $v_1(t), \ldots, v_n(t)$.

The following proposition describes a convenient way to write Eq (1.1). We omit the proof (which is obvious).

PROPOSITION 2.4. *We can rewrite Eq. (1.1) and Eq. (1.2) as follows:*

$$\begin{aligned} y(t+1) &= A(t)v(t) + b(t) \\ q(t+1) &= (1 - \Delta(t))v(t) + \Delta(t)y(t+1) \\ v(t+1) &= q(t+1) + \Delta(t)r(t), \end{aligned}$$

where:

1. *If $i \neq j$ and i, j are neighbors in $G(t)$,*

$$a_{ij}(t) = \frac{1}{4\max(d_i(t), d_j(t))}.$$

However, if $i \neq j$ are not neighbors in $G(t)$, then $a_{ij}(t) = 0$. As a consequence, $A(t)$ is a symmetric matrix.

2. *If node i does not have a measurement of μ at time t, then*

$$a_{ii}(t) = 1 - \frac{1}{4} \sum_{j \in N_i(t), \quad j \neq i} \frac{1}{\max(d_i(t), d_j(t))}.$$

On the other hand, if node i does have a measurement of μ at time t,

$$a_{ii}(t) = \frac{3}{4} - \frac{1}{4} \sum_{j \in N_i(t), \quad j \neq i} \frac{1}{\max(d_i(t), d_j(t))}.$$

Thus $A(t)$ is a diagonally dominant matrix and its graph is merely the intercommunication graph at time t: $G(A(t)) = G(t)$. Moreover, if no node has a measurement at time t, $A(t)$ is stochastic.

3. *If node i does not have a measurement of μ at time t, then $b_i(t) = 0$. If node i does have a measurement of μ at time t, then $b_i(t) = (1/4)\mu$.*
4. *If node i has a measurement of μ at time t, $r_i(t)$ is a random variable with mean zero and variance $\sigma^2/16$. Else, $r_i(t) = 0$. Each $r_i(t)$ is independent of all $v(t)$.*

DEFINITION 2.5. *Let $S(t)$ be the set of nodes with measurement of μ at time t and let $s(t)$ be the cardinality of $S(t)$.*

The following pair of lemmas lower bound the decrease in $Z(v(t))$ from time t to $t+1$. It turns out that we need two distinct lemmas to handle two cases: Lemma 2.6 gives a bound in the case of $s(t)=0$ while Lemma 2.7 handles the case when $s(t)>0$.

LEMMA 2.6. *If $s(t)=0$, then*

$$Z(v(t+1)) \leq Z(v(t)) - \frac{\Delta(t)}{8} \sum_{(k,l)\in E(t)} \frac{(v_i(t)-v_j(t))^2}{\max(d_k(t),d_l(t))}.$$

Proof. By Proposition 2.4, if $s(t)=0$ then $v(t+1)=U(t)v(t)$, where $U(t)=(1-\Delta(t))I+\Delta(t)A$ is a symmetric and stochastic matrix. Therefore

$$v(t+1)-\mu\mathbf{1} = U(t)(v(t)-\mu\mathbf{1}).$$

We apply Corollary 2.2 to obtain

$$||v(t+1)-\mu\mathbf{1}||_2^2 = ||v(t)-\mu\mathbf{1}||_2^2 - \sum_{k<l} [U^2(t)]_{kl}(v_k(t)-v_l(t))^2.$$

We next lower bound $[U^2(t)]_{kl}$ as follows: if $(k,l)\in E(t)$, then $a_{kl}(t)\geq 1/(4\max(d_k(t),d_l(t))$ and consequently $[U]_{kl}(t)\geq \Delta(t)/(4\max(d_k(t),d_l(t))$. Moreover, since $A(t)$ is diagonally dominant so is $U(t)$ and consequently $[U^2]_{kl}\geq \Delta(t)/(8\max(d_i(t),d_j(t)))$. This now immediately implies the statement of the lemma. ∎

LEMMA 2.7. *If $s(t)>0$ and $\Delta(t)\in(0,1)$ then*

$$E[Z(v(t+1) \mid v(t)] \leq Z(v(t)) - \frac{\Delta(t)}{8}\sum_{(k,l)\in E(t)} \frac{(v_k(t)-v_l(t))^2}{\max(d_k(t),d_l(t))} - \frac{\Delta(t)}{4}\sum_{k\in S(t)} (v_k(t)-\mu)^2 + s(t)\frac{\Delta(t)^2}{16}\sigma^2$$
$$\leq \left(1-\frac{1}{8}\Delta(t)\kappa\left[G(t)\right]\right) Z(v(t)) + s(t)\frac{\Delta(t)^2}{16}\sigma^2.$$

Proof. Observe that, for any t, the vector $\mu\mathbf{1}$ satisfies

$$\mu\mathbf{1} = A(t)\mu\mathbf{1} + b(t),$$

and therefore,

$$y(t+1)-\mu\mathbf{1} = A(t)(v(t)-\mu\mathbf{1}). \tag{2.1}$$

We now apply Corollary 2.2 which involves the entries and row-sums of the matrix $A^2(t)$ which we lower-bound as follows. Because $A(t)$ is diagonally dominant and nonnegative, we have that if $(k,l)\in E(t)$ then $[A^2]_{kl}(t)\geq 1/(8\max(d_k(t),d_l(t)))$. Moreover, if k has a measurement of μ then the row sum of the k'th row of A equals $3/4$, which implies that the k'th row sum of A^2 is at most $3/4$. Consequently, Corollary 2.2 implies

$$||y(t+1)-\mu\mathbf{1}||_2^2 \leq Z(v(t)) - \frac{1}{8}\sum_{(k,l)\in E(t)} \frac{(v_k(t)-v_l(t))^2}{\max(d_k(t),d_l(t))} - \frac{1}{4}\sum_{k\in S(t)}(v_k(t)-\mu)^2. \tag{2.2}$$

Next, since $\Delta(t)\in(0,1)$ we can appeal to the convexity of the squared two-norm to obtain

$$||q(t+1)-\mu\mathbf{1}||_2^2 \leq Z(v(t)) - \frac{\Delta(t)}{8}\sum_{k<l,\ (k,l)\in E(t)} \frac{(v_k(t)-v_l(t))^2}{\max(d_k(t),d_l(t))} - \frac{\Delta(t)}{4}\sum_{k\in S(t)}(v_k(t)-\mu)^2.$$

Since $E[r(t)] = 0$ and $E[||r(t)||_2^2] = s(t)\sigma^2/16$ independently of $v(t)$, this immediately implies the first inequality in the statement of the lemma.

To establish the second inequality, we just need to show that

$$\frac{1}{8}\sum_{(k,l)\in E(t)} \frac{(v_k(t)-v_l(t))^2}{\max(d_k(t), d_l(t))} + \frac{1}{4}\sum_{k\in S(t)} (v_k(t)-\mu)^2 \geq \frac{1}{8}\kappa(G(t))Z(v(t)),$$

and the rest of the proof proceeds exactly as after Eq. (2.2). However, the above inequality follows immediately from the definition of $\kappa(G(t))$ after making the change of variables $x_i(t) = v_i(t) - \mu$. □

2.4. Proofs of the main theorems. With the previous pair of lemmas in place, we now go on to provide proofs of Theorems 1-4. Note that it suffices to prove each of these theorems in the case of $k = 1$; once that case is proven, we can apply the theorem to each component of the vector $v_i(t)$ to obtain the general case. Consequently, we will go on assuming that $k = 1$ as in the previous subsection.

Proof. [Proof of Theorem 1.1] We may assume without loss of generality that $\Delta(t) \in (0,1)$ since this is eventually true due to the monotonicity and square summability of $\Delta(t)$. We first claim that there exists some constant $c > 0$ such that if $t_k = 1 + k\max(T, B)$, then

$$E[Z(v(t_{k+1})) \mid v(t_k)] \leq (1 - c\Delta(t_{k+1}))Z(v(t_k)) + n\max(T,B)\Delta(t_k)^2\sigma^2. \qquad (2.3)$$

We postpone the proof of this claim for a few lines while we observe that, as a consequence of our assumptions on $\Delta(t)$, we have the following three facts:

$$\sum_{k=1}^{\infty} c\Delta(t_{k+1}) = +\infty, \qquad \sum_{k=1}^{\infty} n\max(T,B)\Delta(t_k)^2\sigma^2 < \infty, \qquad \lim_{k\to\infty} \frac{n\max(T,B)\Delta(t_k)^2\sigma^2}{c\Delta(t_{k+1})} = 0.$$

Now Lemma 10 from Chapter 2.2 of [25] implies that $\lim_{t\to\infty} Z(v(t)) = 0$ with probability 1.

To conclude the proof, it remains to demonstrate Eq. (2.3). We can apply Lemmas 2.6 and 2.7 to obtain a bound on $E[Z(v(t+1)) \mid v(t)]$ for every time t between t_k and t_{k+1}. Putting these bounds together and taking expectations, we obtain that

$$E[Z(v(t_{k+1})) \mid v(t_k)] \leq Z(v(t_k)) - \sum_{m=t_k}^{t_{k+1}-1} \left(\frac{\Delta(m)}{8} \sum_{(k,l)\in E(m)} \frac{E[(v_k(m)-v_l(m))^2 \mid v(t_k)]}{\max(d_k(m), d_l(m))} \right.$$
$$\left. - \frac{\Delta(m)}{4} \sum_{k\in S(m)} E[(v_k(m)-\mu)^2 \mid v(t_k)] + \Delta^2(m)\frac{\sigma^2}{16}s(m) \right). \quad (2.4)$$

Consequently, Eq. (2.3) follows from the assertion

$$\inf \quad \frac{\sum_{m=t_k}^{t_{k+1}-1} \sum_{(k,l)\in G(m)} E[(v_k(m)-v_l(m))^2 \mid v(t_k)] + \sum_{k\in S(m)} E[(v_k(m)-\mu)^2 \mid v(t_k)]}{\sum_{i=1}^{n} (v_i(t_k)-\mu)^2} > 0$$

where the infimum is taken over all vectors $v(t_k)$ such that $v(t_k) \neq \mu\mathbf{1}$ and over all possible sequence of undirected intercommunication graphs and measurements

between time t_k and $t_{k+1}-1$ satisfying the conditions of uniform connectivity and uniform measurement speed. Now since $E[X^2] \geq E[X]^2$, we have that

$$\inf \quad \frac{\sum_{m=t_k}^{t_{k+1}-1} \sum_{(k,l)\in G(m)} E[(v_k(m)-v_l(m))^2 \mid v(t_k)] + \sum_{k\in S(m)} E[(v_k(m)-\mu)^2 \mid v(t_k)]}{\sum_{i=1}^n (v_i(t_k)-\mu)^2}$$
$$\geq \inf \quad \frac{\sum_{m=t_k}^{t_{k+1}-1} \sum_{(k,l)\in G(m)} E[v_k(m)-v_l(m) \mid v(t_k)]^2 + \sum_{k\in S(m)} E[v_k(m)-\mu \mid v(t_k)]^2}{\sum_{i=1}^n (v_i(t_k)-\mu)^2}.$$

We will complete the proof by arguing that this last infimum is positive.

Let us define $z(t) = E[v(t)-\mu\mathbf{1} \mid v(t_k)]$ for $t \geq t_k$. From Proposition 2.4 and Eq. (2.1), we can work out the dynamics satisfied by the sequence $z(t)$ for $t \geq t_k$:

$$\begin{aligned}
z(t+1) &= E[v(t+1)-\mu\mathbf{1} \mid v(t_k)] \\
&= E[q(t+1)-\mu\mathbf{1} \mid v(t_k)] \\
&= E[(1-\Delta(t))v(t)+\Delta(t)y(t+1)-\mu\mathbf{1} \mid v(t_k)] \\
&= E[(1-\Delta(t))(v(t)-\mu\mathbf{1}) \mid v(t_k)] + E[\Delta(t)(y(t+1)-\mu\mathbf{1}) \mid v(t_k)] \\
&= E[(1-\Delta(t))(v(t)-\mu\mathbf{1}) \mid v(t_k)] + E[\Delta(t)A(t)(v(t)-\mu\mathbf{1}) \mid v(t_k)] \\
&= [(1-\Delta(t))I+\Delta(t)A(t)]\, z(t).
\end{aligned} \tag{2.5}$$

Clearly, we need to argue that

$$\inf \quad \frac{\sum_{m=t_k}^{t_{k+1}-1} \sum_{(k,l)\in G(m)} (z_k(m)-z_l(m))^2 + \sum_{k\in S(m)} z_k^2(m)}{\sum_{i=1}^n z_i^2(t_k)} > 0 \tag{2.6}$$

where the infimum is taken over all sequences of undirected intercommunication graphs satisfying the conditions of uniform connectivity and measurement speed and all nonzero $z(t_k)$ (which in turn determines all the $z(t)$ with $t \geq t_k$ through Eq. (2.5)).

From Eq. (2.5), we have that the expression within in the infimum in Eq. (2.6) is invariant under scaling of $z(t_k)$. So we can conclude that the infimum is achieved by some vector $z(t_k)$ with $||z(t_k)||_2 = 1$.

Now suppose $z(t_k)$ is a vector that achieves this infimum; let $S_+ \subset \{1,\ldots,n\}$ be the set of indices i with $z_i(t_k) > 0$, S_- be the set of indices i with $z_i(t_k) < 0$, and S_0 be the set of indices with $z_i(t_k) = 0$. Since $||z(t_k) - \mu\mathbf{1}||_2 = 1$ we have that at least one of S_+, S_- is nonempty. Without loss of generality, suppose that S_- is nonempty. Due to the conditions of uniform connectivity and uniform measurement speed, there is a first time $t' < t_{k+1}$ when at least one of the following two events happens: (i) some node $i' \in S_-$ is connected to a node $j' \in S_0 \cup S_+$ (ii) some node $i' \in S_-$ has a measurement of μ.

In the former case, $z_{i'}(t') < 0$ and $z_{j'}(t') \geq 0$ and consequently $(z_{i'}(t') - z_{j'}(t'))^2$ will be positive; in the latter case, $z_{i'}(t') < 0$ and consequently $z_{i'}^2(t')$ will be positive. In either case, the infimum of Eq. (2.6) will be strictly positive. ∎

Having proved Theorem 1.1, we next turn to the proof of Theorem 1.2. A crucial part in the proof will be played by the following lemma, which is a modification of a lemma from [11].

Lemma 2.8. *Suppose the sequence b_k satisfies*

$$b_{k+1} \leq (1 - \frac{c_k}{k^{1-\epsilon}})b_k + \frac{d_k}{k^{2-2\epsilon}},$$

where $\epsilon \in (0,1)$ and c_k, d_k are positive sequences and c_k is bounded away from zero. Then

$$\limsup_{k\to\infty} k^{1-\epsilon} b_k \le \sup_k \frac{d_k}{c_k}.$$

Proof. Let $q = \sup_k d_k/c_k$; we may assume that $q < \infty$ because if q is infinite then the lemma is trivially true. Fix $\delta > 0$ and let $q(\delta)$ be a positive number such that $d_k' = c_k q(\delta)$ satisfies $d_k' \ge d_k + \delta$ for all k; such a $q(\delta)$ exists due to the assumptions that $\sup_k d_k/c_k < \infty$ and that c_k is bounded away from zero. Because for fixed $\epsilon \in (0,1)$ we have $1/k^{1-\epsilon} - 1/(k+1)^{1-\epsilon} = O(1/k^{2-\epsilon})$ it follows that for large enough k,

$$\begin{aligned} \frac{d_k}{k^{2-2\epsilon}} &\le \frac{c_k q(\delta)}{k^{2-2\epsilon}} + \frac{q(\delta)}{(k+1)^{1-\epsilon}} - \frac{q(\delta)}{k^{1-\epsilon}} \\ &= \frac{q(\delta)}{(k+1)^{1-\epsilon}} - (1 - \frac{c_k}{k^{1-\epsilon}})\frac{q(\delta)}{k^{1-\epsilon}} \end{aligned}$$

and therefore for large enough k,

$$\begin{aligned} b_k - \frac{q(\delta)}{(k+1)^{1-\epsilon}} &\le (1 - \frac{c_k}{k^{1-\epsilon}})b_k + \frac{d_k}{k^{2-2\epsilon}} - (1 - \frac{c_k}{k^{1-\epsilon}})\frac{q(\delta)}{k^{1-\epsilon}} - \frac{d_k}{k^{2-2\epsilon}} \\ &\le (1 - \frac{c_k}{k^{1-\epsilon}})(b_k - \frac{q(\delta)}{k^{1-\epsilon}}). \end{aligned}$$

Because $\epsilon \in (0,1)$ and c_k is positive and bounded away from zero this implies that $b_k - q(\delta)/k^{1-\epsilon}$ approaches zero faster than the inverse of any polynomial in k. Therefore,

$$\limsup_k k^{1-\epsilon} b_k \le q(\delta).$$

Since this is true for any $\delta > 0$ we have that

$$\limsup_k k^{1-\epsilon} b_k \le q.$$

□

With this lemma in place, we now proceed to the proof of Theorem 1.2. The proof has many parallels to the proof of Theorem 1.1, in particular in that it relies on the repeated application of Lemmas 2.6 and 2.7 to produce a recursion which bounds $E[Z(v(t)) \mid v(t_k)]$. Lemma 2.8, which we have just proved, will be used to obtain a precise convergence time decay out of this recursion.

Proof. [Proof of Theorem 1.2] Let t_k be one plus the time when the k'th measurement occurs; naturally, $t_k \le 1 + kT$. Lemmas 2.6 and 2.7 imply that

$$E[Z(v(t_{k+1})) \mid v(t_k)] \le (1 - \frac{1}{8}\Delta(t_{k+1})\kappa[G(t_{k+1})])Z(v(t_k)) + s(t_{k+1})\Delta(t_{k+1})^2\frac{\sigma^2}{16}.$$

Tautologically $\kappa(G(t)) \ge \min_t \kappa(G(t))$, and therefore

$$E[Z(v(t_{k+1})) \mid v(t_k)] \le (1 - \frac{1}{8}\Delta(t_{k+1})\min_t \kappa\,[G(t)])Z(v(t_k)) + n\Delta(t_{k+1})^2\frac{\sigma^2}{16}.$$

Plugging in $\Delta(t) = 1/t^{1-\epsilon}$, we obtain

$$E[Z(v(t_{k+1})) \mid v(t_k)] \leq (1 - \frac{(1/8)\min_t \kappa\,[G(t)]}{t_{k+1}^{1-\epsilon}})Z(v(t_k)) + \frac{n\sigma^2/16}{t_{k+1}^{2-2\epsilon}}.$$

Let $\overline{t}_k = k/t_{k+1}$. We may rewrite the above relation as

$$E[Z(v(t_{k+1})) \mid v(t_k)] \leq (1 - \frac{\overline{t}_k^{1-\epsilon}(1/8)\min_t \kappa\,[G(t)]}{k^{1-\epsilon}})Z(v(t_k)) + \frac{\overline{t}_k^{2-2\epsilon} n\sigma^2/16}{k^{2-2\epsilon}}.$$

Taking expectations of both sides,

$$E[Z(v(t_{k+1})) \mid v(0)] \leq (1 - \frac{\overline{t}_k^{1-\epsilon}(1/8)\min_t \kappa\,[G(t)]}{k^{1-\epsilon}})E[Z(v(t_k)) \mid v(0)] + \frac{\overline{t}_k^{2-2\epsilon} n\sigma^2/16}{k^{2-2\epsilon}}.$$

We now argue that the assumptions of Lemma 2.8 apply. All that needs to be verified is that $\overline{t}_k$ is bounded above and away from zero, which is, of course, true: $1 \geq \overline{t}_k \geq 1/(2T)$ because $k \leq t_{k+1} \leq (k+1)T$. Application of Lemma 2.8 gives

$$\limsup_k k^{1-\epsilon} E[Z(v(t_k)) \mid v(0)] \leq \frac{n\sigma^2}{2\min_t \kappa\,[G(t)]}. \tag{2.7}$$

Finally, to complete the proof we must pass from a statement about the decay of $E[Z(v(t_k)) \mid v(0)]$ at the instances t_k to a statement about the decay of $E[Z(v(t)) \mid v(0)]$ for general t. So suppose that prior to time t there have been $m(t)$ instants when a node has had a measurement. Because $E[Z(v(t+1)) \mid v(t)] \leq Z(v(t))$, if t is not among the t_k then Eq. (2.7) implies

$$\limsup_t (m(t))^{1-\epsilon} E[Z(v(t)) \mid v(0)] \leq \frac{n\sigma^2}{2\min_t \kappa\,[G(t)]}.$$

Since $m(t) \geq (t-1)/T$, the last equation implies Theorem 1.2. ☐

We now turn to the proof of Theorem 1.4. Its proof requires a certain inequality between quadratic forms in the vector $v(t)$ which we separate into the following lemma.

Lemma 2.9. *Let $t_k = 1 + k\max(T, B)$ and assume that the entries of the vector $v(t_k)$ are ordered monotonically as*

$$v_1(t_k) < v_2(t_k) < \cdots < v_n(t_k).$$

Further, let us assume that none of the $v_i(t_k)$ equal μ, and let us define p_- to be the largest index such that $v_{p-}(t_k) < \mu$ and p_+ to be the smallest index such that $v_p(t_k) > \mu$ is nonnegative. We then have

$$\sum_{m=t_k}^{t_{k+1}-1} \sum_{(k,l)\in E(m)} E[(v_k(m) - v_l(m))^2 \mid v(t_k)] + \sum_{k\in S(m)} E[(v_k(m) - \mu)^2 \mid v(t_k)] \geq$$
$$(v_{p_-}(t_k) - \mu)^2 + (v_{p+}(t_k) - \mu)^2 + \sum_{i=1,\ldots,n,\quad i\neq p_-} (v_i(t_k) - v_{i+1}(t_k))^2. \tag{2.8}$$

Proof. The proof parallels a portion of the proof of Theorem 1.1. First, we change variables by defining $z(t)$ as

$$z(t) = E[z(t) - \mu\mathbf{1} \mid v(t_k)]$$

for $t \geq t_k$. We claim that

$$\sum_{m=t_k}^{t_{k+1}-1} \sum_{(k,l)\in E(m)} (z_k(m)-z_l(m))^2 + \sum_{k\in S(m)} z_k^2(m) \geq z_{p_-}^2(t_k)+z_{p_+}^2(t_k)+\sum_{i=1}^{n-1}(z_i(t_k)-z_{i+1}(t_k))^2 \tag{2.9}$$

The claim immediately implies the lemma after application of the inequality $E[X^2] \geq E[X]^2$.

Now we turn to the proof of the claim, which is similar to a proof of lemma from [22]. We will associate to each term on the right-hand side of Eq. (2.9) a term on the left-hand side of Eq. (2.9), and we will argue that each term on the left-hand side is at least is big as the sum of all terms on the right-hand side associated with it. Since every term on the left-hand side is clearly nonnegative, this will prove the lemma.

To describe this association, we first introduce some new notation. We denote the set $\{1, \ldots, l\}$ by S_l; its complement the set $\{l+1, \ldots, n\}$ is then S_l^c. If $l \neq p_-$, we will abuse notation by saying that S_l contains zero if $l \geq p_+$; else, we say that S_l does not contain zero and S_l^c contains zero. However, in the case of $l = p_-$, we will say that neither of S_{p_-} and $S_{p_-}^c$ contains zero. We will say that S_l "is crossed by an edge" at time m if a node in S_l is connected to a node in S_l^c at time m. For $l \neq p_-$, we will say that S_l is "crossed by a measurement" at time m if a node in whichever of S_l, S_l^c that does not contain zero has a measurement at time m. We will say that S_{p_-} is "crossed by a measurement from the left" at time m if a node in S_{p_-} has a measurement at time m; we will say that it is "crossed by a measurement from the right" at time m if a node in $S_{p_-}^c$ had a measurement at time m. Note that the assumption of uniform connectivity means that every S_l is crossed by an edge at least one time $m \in \{t_k, \ldots, t_{k+1}-1\}$. It may happen that some S_l are also crossed by measurements, but it isn't required by the uniform measurement assumption. Nevertheless, the uniform measurement assumption implies that S_{p_-} is crossed by a measurement at some time $m \in \{t_k, \ldots, t_{k+1}-1\}$. Finally, we will say that S_l is crossed at time m if it is either crossed by an edge or crossed by a measurement (plainly or from left or right).

We next describe how we associate terms on the right-hand side of Eq. (2.9) with terms on the left-hand side of Eq. (2.9). Suppose l is any number in $1, \ldots, n-1$ except p_-; consider the *first* time S_l is crossed; let this be time m. If the crossing is by an edge, then let (i,j) be any edge which goes between S_l and S_l^c at time m. We will associate $(z_l(t_k) - z_{l+1}(t_k))^2$ with $(z_i(m) - z_j(m))^2$; as a shorthand for this, we will say that we associate l with the edge (i,j) at time m. On the other hand, if S_l is crossed by a measurement[2] at time m, let i be a node in whichever of S_l, S_l^c does not contain zero which has a measurement at time m; we associate $(z_l(t_k)-z_{l+1}(t_k))^2$ with $z_i^2(m)$; as a shorthand for this, we will say that we associate l with a measurement by i at time m.

Finally, we describe the associations for the terms $v_{p_-}(t_k)^2$ and $v_{p+}(t_k)^2$, which are more intricate. Again, let us suppose that that S_{p_-} is crossed first at time m; if the crossing is by an edge, then we associate both these terms with any edge (i,j) crossing S_{p_-} at time m. If, however, S_{p_-} is crossed first by a measurement from the left, then we associate $v_{p_-}^2(t_k)$ with $z_i^2(m)$, where i is any node in S_{p_-} having a

[2]If S_l is crossed both by an edge and a measurement at time m, we will say it is crossed by an edge first. Throughout the remainder of this proof, we keep to the convention breaking ties in favor of edges by saying that S_l is crossed first by an edge if the first crossing was simultaneously by both an edge and by a measurement.

measurement at time m. We then consider u, which is the first time S_{p_-} is crossed by either an edge or a measurement from the right; if it is crossed by an edge, then we associate $v_{p_+}(t_k)$ with $(z_i(u) - z_j(u))^2$ with any edge (i,j) at going between S_{p_-} and $S_{p_-}^c$ at time u; else, we associate it with $z_i^2(u)$ where i is any node in $S_{p_-}^c$ having a measurement at time u. On the other hand, if S_{p_-} is crossed first by a measurement from the right, then we flip the associations: we associate $v_{p_+}^2(t_k)$ with $z_i^2(m)$, where i is any node in $S_{p_-}^c$ having a measurement at time m. We then consider u, which is now the first time S_{p_-} is crossed by either an edge or a measurement from the left; if S_{p_-} is crossed by an edge first, then we associate $v_{p-}(t_k)$ with $(z_i(u) - z_j(u))^2$ with any edge (i,j) at going between S_{p_-} and $S_{p_-}^c$ at time u; else, we associate it with $z_i^2(u)$ where i is any node in S_{p_-} having a measurement at time u.

We now go on to prove that every term on the left-hand side of Eq. (2.9) is at least as big as the sum of all terms on the right-hand side of Eq. (2.9) associated with it.

Let us first consider the terms $(z_i(m) - z_j(m))^2$ on the left-hand side of Eq. (2.9). Suppose the edge (i,j) with $i < j$ at time m was associated with indices $l_1 < l_2 < \cdots < l_r$ (i.e., with the terms $(z_{l_1}(t_k) - z_{l_1+1}(t_k))^2, \ldots, (z_{l_r}(t_k) - z_{l_r+1}(t_k))^2$. The key observation is that if S_l has not been crossed before time m then

$$\max_{i=1,\ldots,l} z_i(m) \le z_l(t_k) \le z_{l+1}(t_k) \le \min_{i=l+1,\ldots,n} z_i(m).$$

Consequently,

$$z_i(m) \le z_{l_1}(t_k) \le z_{l_1+1}(t_k) \le z_{l_2}(t_k) \le z_{l_2+1}(t_k) \le \cdots \le z_{l_r}(t_k) \le z_{l_r+1}(t_k) \le z_j(m)$$

which implies that

$$(z_i(m)-z_j(m))^2 \ge (z_{l_1+1}(t_k)-z_{l_1}(t_k))^2+(z_{l_2+1}(t_k)-z_{l_2}(t_k))^2+\cdots+(z_{l_r+1}(t_k)-z_{l_r}(t_k))^2.$$

This proves the statement in the case when the edge (i,j) is associated with $l_1 < l_2 < \cdots < l_r$.

Suppose now that the edge (i,j) is associated with indices $l_1 < l_2 < \cdots < l_r$ as well as both the terms $z_{p_-}^2(t_k), z_{p_+}^2(t_k)$. This happens when every S_{l_i} and S_{p_-} is crossed for the first time by (i,j), so that we can simply repeat the sequence of steps in the previous paragraph to obtain

$$(z_i(m)-z_j(m))^2 \ge (z_{l_1+1}(t_k)-z_{l_1}(t_k))^2+(z_{l_2+1}(t_k)-z_{l_2}(t_k))^2+\cdots+(z_{l_r+1}(t_k)-z_{l_r}(t_k))^2+(z_{p_-}(t_k)+z_{p_+}(t_k))^2$$

which, since $(z_{p_-}(t_k)-z_{p_+}(t_k))^2 \ge z_{p_-}^2(t_k)+z_{p_+}^2(t_k)$ proves the statement in this case.

Suppose now that the edge (i,j) with $i < j$ at time m is associated with indices $l_1 < l_2 < \cdots < l_r$ as well as the term $z_{p_-}^2(t_k)$. This happens when every l_i has not been crossed before time m, S_{p_-} is being crossed by an edge at time m and has been crossed from the right but not from the left before time m. Consequently, in addition to the inequalities $i \le l_1, j \ge l_r + 1$ we have the additional inequalities $i \le p_-$ while $j \ge p_+$ (since (i,j) crosses S_{p_-}). Because S_{p_-} has not been crossed by an edge before, we have $z_j(m) > 0$, so that

$$(z_i(m)-z_j(m))^2 \ge (z_{l_1+1}(t_k)-z_{l_1}(t_k))^2+(z_{l_2+1}(t_k)-z_{l_2}(t_k))^2+\cdots+(z_{l_r+1}(t_k)-z_{l_r}(t_k))^2+(z_{p_-}(t_k)-0)^2(t_k)$$

which proves the statement in this case.

The proof when the edge (i,j) is associated with index $l_1 < \cdots < l_r$ and $z_{p_+}^2(t_k)$ is similar, and we omit it. Consequently, we have now proved the desired statement for all the terms of the form $(z_i(m) - z_j(m))^2$.

It remains to consider the terms $z_i^2(m)$. So let us suppose that the term $z_i^2(m)$ is associated with indices $l_1 < l_2 < \cdots < l_r$ as well as possibly one of $z_{p_-}^2(t_k), z_{p_+}^2(t_k)$ (due to the way we defined the associations, it can never be associated with both). Observe that we must have either $i \leq l_1$ (if S_{l_r} does not contain zero) or $i > l_r$ (if S_{l_r} contains zero). We suppose that it is the former case; the proof in the latter case is similar. Thus we must have $i \leq l_r < p_-$ so it is not possible for $z_i^2(m)$ to be associated with $z_{p_+}^2(t_k)$; however, $v_{p_-}^2(t_k)$ might still be associated with it. Since S_{l_1} has not been crossed before, we have that

$$z_i(m) \leq z_{l_1}(t_k) \leq z_{l_1+1}(t_k) \leq z_{l_2}(t_k) \leq z_{l_2+1}(t_k) \leq \cdots \leq z_{l_r}(t_k) \leq z_{l_r+1}(t_k) \leq z_{p_-}(t_k) < 0$$

and therefore

$$z_i^2(m) \geq (z_{l_1+1}(t_k) - z_{l_1}(t_k))^2 + (z_{l_2+1}(t_k) - z_{l_2}(t_k))^2 + \cdots + (z_{l_r+1}(t_k) - z_{l_r}(t_k))^2 + (z_{p_-}(t_k) - 0)^2$$

which concludes the proof. □

We now put all the pieces together and provide a proof of Theorem 1.4.

Proof. [Proof of Theorem 1.4] As in the statement of Lemma 2.9, let us choose $t_k = 1 + k\max(T, B)$. Observe that by continuity Lemma (2.8) holds even with the strict inequalities between $v_i(t_k)$ replaced with nonstrict inequalities and without the assumption that none of the $v_i(t_k)$ are zero; moreover, using the inequality

$$(v_{p_-}(t_k) - \mu)^2 + (v_{p+}(t_k) - \mu)^2 \geq \frac{(v_{p_-}(t_k) - \mu)^2 + (v_{p+}(t_k) - \mu)^2 + (v_{p_-}(t_k) - v_{p_+}(t_k))^2}{4},$$

we have that Lemma (2.8) implies that

$$\sum_{m=t_k}^{t_{k+1}-1} \sum_{(k,l) \in E(m)} E[(v_k(m) - v_l(m))^2 \mid v(t_k)] + \sum_{k \in S(m)} E[(v_k(m) - \mu)^2 \mid v(t_k)] \geq \frac{1}{4}\kappa(L_n) Z(v(t_k)).$$

Because $\Delta(t)$ is decreasing and the degree of any vertex at any time is at most $d_{\max}$, this in turn implies

$$\sum_{m=t_k}^{t_{k+1}-1} \frac{\Delta(m)}{8} \sum_{(k,l) \in E(m)} \frac{E[(v_k(m) - v_l(m))^2 \mid v(t_k)]}{\max(d_k(m), d_l(m))} + \frac{\Delta(m)}{4} \sum_{k \in S(m)} E[(v_k(m) - \mu)^2 \mid v(t_k)] \geq \frac{\Delta(t_{k+1})}{32 d_{\max}} \kappa(L_n) Z(v(t_k)).$$

Now appealing to Eq. (2.4), we have

$$E[Z(v(t_{k+1})) \mid v(t_k)] \leq \left(1 - \frac{\Delta(t_{k+1})\kappa(L_n)}{32 d_{\max}}\right) Z(v(t_k)) + n\Delta(t_k)^2 \max(T, B) \frac{\sigma^2}{16}.$$

Taking expectations and using Lemma 2.8,

$$\limsup_{t_k} \left(\frac{t_k - 1}{\max(T, B)}\right)^{1-\epsilon} E[Z(v(t_k)) \mid v(0)] \leq \frac{2 n d_{\max} \max(T, B) \sigma^2}{\kappa(L_n)},$$

which implies

$$\limsup_{t_k} t_k^{1-\epsilon} E[Z(v(t_k)) \mid v(0)] \leq \frac{2 n d_{\max} \max(T, B)^{2-\epsilon} \sigma^2}{\kappa(L_n)}. \tag{2.10}$$

This is nearly the statement of the theorem; we need an argument which will allow us to replace t_k by an arbitrary t. We next argue that this is possible at the expense of adding a factor of 2 to the right-hand side. Indeed, Corollary 2.2 implies that $||Qx||_2 \leq ||x||_2$ for any symmetric nonnegative stochastic matrix Q. Consequently for any t between t_k and t_{k+1}, we have

$$E[Z(v(t+1)) \mid v(t)] \leq Z(v(t)) + n\Delta^2(t)\sigma^2/16.$$

Taking expectations,

$$E[Z(v(t)) \mid v(t_k)] \leq Z(v(t_k)) + n\Delta^2(t_k)\max(T,B)\sigma^2/16.$$

We combine this inequality with Eq. (2.10) in the following chain of inequalities: letting p_k be the largest t_k which is at most t, and observing that for $t \geq 2\max(T,B)+1$ we have $t \leq 2p_k$ for $k \geq 1$, we have

$$\begin{aligned}
\limsup_t t^{1-\epsilon} E[Z(v(t)) \mid v(0)] &\leq \limsup_t t^{1-\epsilon} E[E[Z(v(t) \mid v(p_k), v(0)]] \\
&\leq \limsup_t t^{1-\epsilon} \left(E[Z(v(p_k)) \mid v(0)] + n\Delta^2(p_k)\max(T,B)\sigma^2/16 \right) \\
&\leq \limsup_k 2p_k^{1-\epsilon} \left(E[Z(v(p_k)) \mid v(0)] + n\frac{1}{p_k^{2-2\epsilon}}\max(T,B)\sigma^2/16 \right) \\
&\leq \frac{8nd_{\max}\max(T,B)^{2-\epsilon}\sigma^2}{\kappa(L_n)} + 0.
\end{aligned}$$

The theorem is now proved.
□

We finally turn to the proof of our last major result, Theorem 1.3. We note that this proof has close parallels with the arguments used in [16] to prove eigenvalue bounds for stochastic matrices.

Proof. [Proof of Theorem 1.3] We will show that for any m,

$$\min_{||x||_2=1} \qquad x_m^2 + \sum_{(i,j)\in E}(x_i - x_j)^2 \geq \frac{1}{D_{\max} n}$$

This then implies the theorem immediately from the definition of the sieve constant.

Indeed, we may suppose $m = 1$ without loss of generality. Suppose the minimum in the above optimization problem is achieved by the vector x; let M be the index of the component of x with the largest absolute value; without loss of generality, we may suppose that the shortest path connecting 1 and M is $1 - 2 - \cdots - M$ (we can simply relabel the nodes to make this true). Moreover, we may also assume $x_M > 0$ (else, we can just replace x with $-x$).

Now the assumptions that $||x||_2 = 1$, that x_M is the largest component of x in absolute value, and that $x_M > 0$ imply that $x_M \geq 1/\sqrt{n}$ or

$$(x_1 - 0) + (x_2 - x_1) + \cdots + (x_M - x_{M-1}) \geq \frac{1}{\sqrt{n}}$$

and applying Cauchy-Schwarz

$$M(x_1^2 + (x_2 - x_1)^2 + \cdots (x_M - x_{M-1})^2) \geq \frac{1}{n},$$

or

$$x_1^2 + (x_2 - x_1)^2 + \cdots (x_M - x_{M-1})^2 \geq \frac{1}{Mn} \geq \frac{1}{D_{\max} n}.$$

$\square$

We end this section with an extended remark discussing how our results may be ported to prove guarantees for cooperative learning in settings when nodes unpredictably enter and leave the system.

Remark 2.10. *We remark that our learning protocol is capable of coping with persistent node failures. Consider, for example, the scenario in which nodes sometimes "terminate," i.e., nodes are occasionally removed permanently from the system. Note that the proofs of our main theorems work by bounding how much $Z(v(t))$ decreases and the termination of a node can only decrease $Z(v(t))$. Thus our proofs can be easily extended to cover this possibility. One can assert, for example, that convergence to μ happens with probability 1 to μ as long as at least one node does not terminate. In fact, since termination only decreases n, we can further assert that all of the Theorems 1-4 immediately carry over to the setting when nodes are allowed to terminate with only the additional proviso that at least one node does not terminate.*

Conversely, suppose that new nodes are allowed to periodically join the system. The decay bounds that we have derived can be applied immediately after a new node joins. We can assert, for example, that as long as finitely many new nodes join the system, convergence to μ occurs with probability 1. To derive quantitative convergence bounds, we need some control over the values of new nodes that join the system; if these are far from μ, convergence may take an arbitrarily long time.

It seems natural to suppose that new nodes join the system with a value which lies in $[\mu - L, \mu + U]$, where L and U are some upper and lower bounds. Consequently, the addition of a new node to the system increases $Z(v(t))$ by at most $\max(L, U)^2$, after which our decay bounds of Lemmas 2.6 and 2.7 immediately apply. Thus depending on the time that the joins occur, we may derive convergence rate bounds. For example, if we know that j nodes join between times 1 and k, but none after, a version of Theorem 3 can be proven with the following conclusion:

$$\limsup_{t\to\infty} (t-k)^{1-\epsilon} E[Z(v(t)) \mid v(0)] \leq \frac{4nd_{\max}\sigma^2 \max(T,B)^{2-\epsilon}(Z(0) + j\max(L,U)^2 + k(n+j)\sigma^2/16)}{\kappa(L_n)}.$$

The case when both joins and terminations are allowed is not straightforward; it is clear that convergence to μ may be impossible if, say, a node terminates after every time it measures μ and before it contacts other nodes. However, we can nevertheless prove analogues of Theorems 1-4 in a variety of settings.

For example, suppose that at least one node has a measurement every T steps and does not terminate in the following B steps. Suppose further that nodes which join the network do so with values that are equal to a value of one of the nodes in the network. This implies that a join at most doubles $Z(v(t))$. Consequently, as long as joins happen less frequently with time - say, slower than the bounds of Theorems 2 and 3 imply that $E[Z(v(t))]$ gets multiplied by 1/4 - convergence with probability 1 to μ still occurs.

3. The sieve constant. The previous section showed that the scaling of our learning protocol with network structure can be upper bounded in terms of the sieve constant. In this section, we proceed to compute sieve constants for various common

graphs. The bounds we derive in this section may be immediately translated into improved bounds on the performance of our learning protocol on certain classes of graphs. For example, we show that for the complete graph K_n, the sieve constant satisfies the upper bound $1/\kappa(K_n) \leq cn$ for some constant c; this is two orders of magnitude better than the bound $1/\kappa(G_n) \leq n^3$ which follows from Theorem 1.3, and it may be plugged in directly into the statement of Theorem 1.2 to get an improved convergence for our protocol on the complete graph.

We are primarily interested in the scaling of the sieve constant with the number of nodes n, and correspondingly we will be satisfied to compute sieve constants to within a constant factor. We will extensively use the notations $f(n) = \Omega(g(n))$ and $f(n) = \Theta(g(n))$ which mean, respectively, $f(n) \geq cg(n)$, and $cg(n) \leq f(n) \leq Cg(n)$, for some constants c, C not depending on n. We begin with a lemma which will simplify some of the forthcoming computations.

LEMMA 3.1. *Let $G = (V, E)$ be an undirected, connected graph and let $\mathcal{P}$ be the group of automorphisms of G which fix the vertex m. If the minimization problem*

$$\min_{x \in \mathbb{R}^n} \frac{x_m^2 + \sum_{(i,j) \in E} (1/\max(d_i, d_j))(x_i - x_j)^2}{\sum_{i=1}^n x_i^2}$$

has an optimal solution with $x_m \neq 0$ then it has a solution that is invariant under the actions of any $P \in \mathcal{P}$.

Proof. We make a few preliminary remarks before turning to the proof. Let us define E_{ij} to be the matrix with 1 in the (i,j)'th place and zero in all other entries. By the Courant-Fischer theorem, the optimal value in the above optimization problem is the smallest eigenvalue λ_n of the matrix $B = E_{mm} + L$, where L is the Laplacian matrix of the graph G with edge weights $w_{ij} = 1/(\max(d_i, d_j))$. Moreover, the set of vectors achieving the optimal value is precisely the set of all nonzero vectors in the eigenspace corresponding to that eigenvalue.

The condition that $\mathcal{P}$ is the group of automorphisms of G which preserve m may be stated as follows: $P \in \mathcal{P}$ if and only if P is a permutation matrix satisfying $Pe_m = e_m$ and $P^{-1}BP = B$. Therefore, if x is an eigenvector of B with eigenvalue λ, then Px is also an eigenvector of B with eigenvalue λ:

$$BPx = PP^{-1}BPx = PBx = \lambda Px.$$

We turn now to the proof of the lemma. Suppose x is a vector which achieves the minimum in the above minimization problem. For any $P \in \mathcal{P}$, let $i(P)$ be the smallest integer so that $P^{i(P)}$ is the identity permutation. Define

$$y = \sum_{P \in \mathcal{P}} \quad \sum_{j=1,\ldots,i(P)} \quad P^j x$$

Clearly, y is invariant under the action of any $P \in \mathcal{P}$. Because x is an eigenvector of B corresponding to the smallest eigenvalue, so is y; moreover, $y \neq 0$ because $x_m \neq 0$ and every power of $P \in \mathcal{P}$ fixes m. Consequently, y achieves the minimum of the optimization problem. This concludes the proof. ∎

The following series of propositions gives the sieve constants of some common graphs up to a constant factor.

PROPOSITION 3.2. *The sieve constant of the complete graph K_n is $\Theta(1/n)$.*

Proof. As a consequence of Lemma 3.1 we have two cases to consider. In the first case $x_m = 0$ and then the value of $\kappa(K_n)$ is equal to the optimal value of the minimization problem

$$\kappa(K_n) = \min_{||x||_2=1, x_m=0} \quad \frac{1}{n-1} \sum_{k \neq m, l \neq m} (x_k - x_l)^2 + \frac{1}{n-1} \sum_{i=1,\ldots,n, i \neq m} x_i^2.$$

Since the second term always equals $1/(n-1)$, the best we can do is pick an x that sets the first term equal to zero; this yields the value $1/(n-1)$ for the optimal value of the minimization program.

In the second case, $x_m \neq 0$ but then by Lemma 3.1 the value of every x_i with $i \neq m$ is the same; consequently, in this case the value of $\kappa(K_n)$ is equal to the optimal value of the minimization problem

$$\kappa(K_n) = \min_{a^2+(n-1)b^2=1} a^2 + (b-a)^2.$$

Observe that if $a \geq 1/\sqrt{3n}$, then the objective function is at least $1/(3n)$; and if $a < 1/\sqrt{3n}$ then $(n-1)b^2 \geq (3n-1)/(3n)$ so $b^2 \geq 1/n$; in turn, this means the objective function is at least $(1/\sqrt{n} - 1/\sqrt{3n})^2$ which is $\Omega(1/n)$. Since we can find an x which achieves an objective value of $O(1/n)$ (as we saw in the previous paragraph) we can conclude that $\kappa(K_n) = \Theta(1/n)$. ∎

PROPOSITION 3.3. *The sieve constant of two complete graphs connected by an edge (see Figure 3.1), which we denote $K_n - K_n$, is $\Theta(1/n^2)$.*

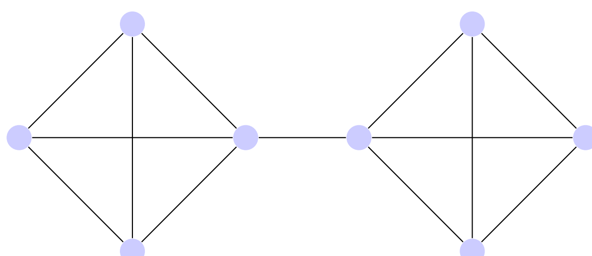

FIG. 3.1. *The graph $K_n - K_n$ for $n = 4$.*

Proof. By Theorem 1.3, $\kappa(K_n - K_n) \geq 1/(3n^2)$. On the other hand, let m be any point in the first complete graph, and set $x_i = 0$ for all i in the first complete graph and $x_i = 1/\sqrt{n/2}$ for all i in the second. This yields

$$\kappa(K_n - K_n) \leq \frac{1}{n/2+1} \frac{2}{n}$$

so that $\kappa(K_n - K_n) = \Theta(1/n^2)$. ∎

PROPOSITION 3.4. *The sieve constant of the line graph L_n is $\Theta(1/n^2)$.*

Proof. Observe that by Theorem 1.3, $\kappa(L_n) \geq 1/(2n^2)$. On the other hand, number the vertices of the line $1, \ldots, m$ from left to right, set $m = 1$ and pick $x_i = i/n^{1.5}$. Then

$$\kappa(L_n) \leq \frac{1/n^3 + \sum_{i=1}^n 1/n^3}{\sum_{i=1}^n i^2/n^3} \leq O\left(\frac{1}{n^2}\right).$$

Consequently, $\kappa(L_n) = \Theta(1/n^2)$. ∎

PROPOSITION 3.5. *The sieve constant of the ring graph R_n is $\Theta(1/n^2)$.*

Proof. Theorem 1.3 gives $\kappa(R_n) \geq 1/(2n^2)$. Suppose n is even. Number the nodes $1, \ldots, m$ counterclockwise and pick $m = 1$. Setting

$$\begin{cases} x_i = i/n^{1.5} & \text{if } i = 1, \ldots, n/2 \\ x_i = \frac{n-i}{n^{1.5}} & \text{if } i = n/2+1, \ldots, n \end{cases}$$

yields the upper bound

$$\kappa(R_n) \leq \frac{1/n^3 + n(1/n^3)}{\sum_{i=1}^{n/2} i^2/n^3 + \sum_{i=0}^{n/2-1} i^2/n^3} \leq O\left(\frac{1}{n^2}\right).$$

Consequently, $\kappa(L_n) = \Theta(1/n^2)$. A similar argument proves the same conclusion when n is odd. ∎

PROPOSITION 3.6. *The sieve constant of the star graph S_n (see Figure 3.2) is $\Theta(1/n^2)$.*

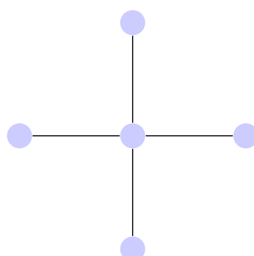

FIG. 3.2. *The star graph S_4.*

Proof. Theorem 1.3 gives that $\kappa(S_n) \geq 1/(2n^2)$ because the diameter is 2. On the other hand, choose m to be a leaf and set $x_m = 0$ and $x_i = 1/\sqrt{n-1}$ elsewhere. This yields

$$\kappa(S_n) \leq \frac{1}{(n-1)^2},$$

which proves this proposition. ∎

4. Simulations. We report here on several simulations of our learning protocol. These simulations confirm the broad outlines of the bounds we have derived; the convergence to μ takes place at a rate broadly consistent with a decay in $1/t$ and the scaling with n appears to be polynomial.

Figure 4.1 shows plots of the distance from μ for the complete graph, the line graph (with one of the endpoint nodes doing the sampling), and the star graph (with the center node doing the sampling), each on 40 nodes. We caution that there is no reason to believe these charts capture the correct asymptotic behavior as $t \to \infty$. Nevertheless, based on the performance shown in these plots, we see that a linear decay in the distance from μ with time appears to be plausible.

Intriguingly, the star graph and the complete graph appear to have very similar performance. By contrast, the performance of the line graph is an order of magnitude

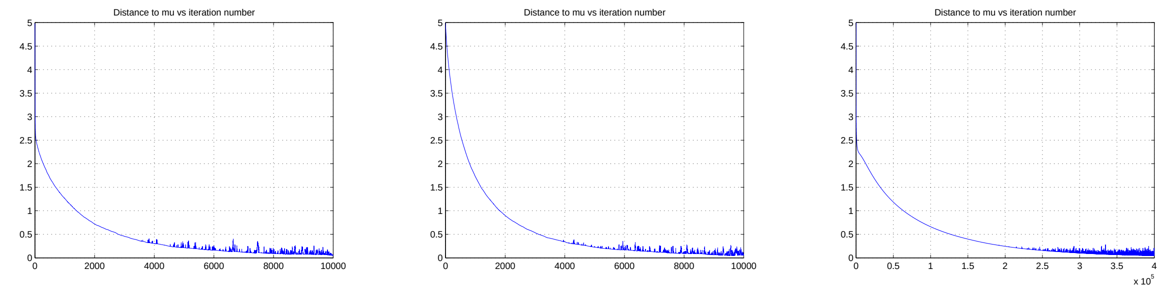

FIG. 4.1. *The three plots show the quantity $||v(t) - \mu\mathbf{1}||_\infty$ as a function of the number of iterations. The graph on the left corresponds to the complete graph, the middle graph corresponds to the star graph, and the graph on the right corresponds to the line graph. Each graph has* 40 *vertices and in each case, exactly one node is doing the measurements; in the star graph it is the center vertex and in the line graph it is one of the endpoint vertices. The initial vector is random with entries uniform in* $[0, 5]$ *in each case. Stepsize* Δ *is chosen to be* $1/t^{1/4}$ *for all three simulations.*

inferior to the performance of either of these; it takes the line graph on 40 nodes on the order of $400,000$ iterations to reach roughly the same level of accuracy that the complete graph and star graph reach after about $10,000$ iterations.

Next, Figure 4.2 focuses on the scaling with the number of nodes n. The graphs show the time until $||v(t) - \mu\mathbf{1}||_\infty$ decreases below a certain threshold as a function of number of nodes. We see scaling that could plausibly be quadratic for the line graph and linear for the complete graph, which matches the upper bounds we have derived in Section 3. However, we see scaling which appears linear for the star graph which is an order-of-magnitude better than the quadratic upper bound of Section 3.

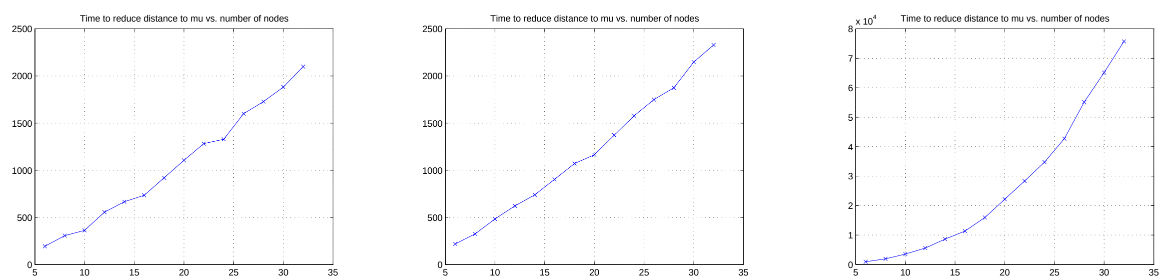

FIG. 4.2. *The three plots show how long it takes the quantity $||v(t) - \mu\mathbf{1}||_\infty$ to shrink below* $1/2$ *starting from a random vector with entries in* $[0, 5]$*. The graph on the left corresponds to the complete graph, the middle graph corresponds to the star graph, and the graph on the right corresponds to the line graph. In each case, exactly one node is doing the measurements; in the star graph it is the center vertex and in the line graph it is one of the endpoint vertices. Stepsize is chosen to be* $1/t^{1/4}$ *for all three simulations.*

Finally, we include a simulation for the lollipop graph, defined to be a complete graph on $n/2$ vertices joined to a line graph on $n/2$ vertices. The lollipop graph often appears as an extremal graph for various random walk properties (see, for example, [8]). The node at the end of the stem, i.e., the node which is furthest from the complete subgraph, is doing the sampling. The scaling with the number of nodes is considerably worse than for the other graphs we have simulated here.

Finally, we emphasize that the learning speed also depends on the precise location

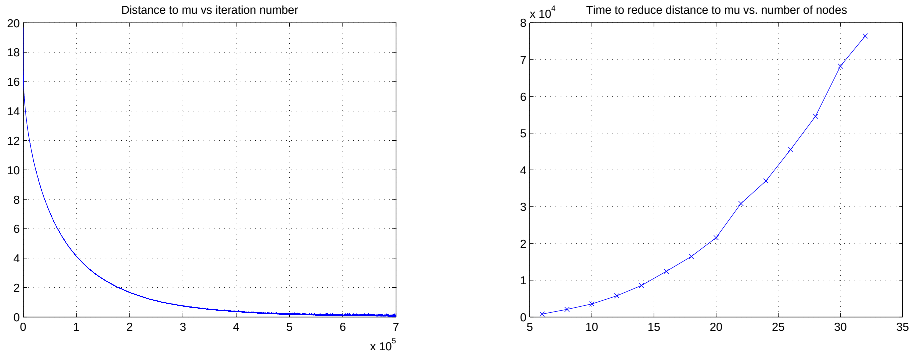

FIG. 4.3. *The plot on the left shows $||v(t) - \mu\mathbf{1}||_\infty$ as a function of the number of iterations for the lollipop graph on 40 nodes; the plot on the right shows the time until $||v(t) - \mu\mathbf{1}||_\infty$ shrinks below 0.5 as function of the number of nodes n. In each case, exactly one node is performing the measurements, and it is the node farthest from the complete subgraph. The starting point is a random vector with entires in $[0, 5]$ for both simulation and stepsize is $1/t^{1/4}$.*

of the node doing the sampling within the graph. While our results in this paper bound the worst case performance over all choices of sampling node, it may very well be that by appropriately choosing the sensing nodes, better performance relative to our bounds and relative to these simulations can be achieved.

5. Conclusion. We have proposed a model for cooperative learning by multi-agent systems facing time-varying connectivity and intermittent measurements. We have proved a protocol capable of learning an unknown vector from independent measurements in this setting and provided quantitative bounds on its learning speed. Crucially, these bounds have a dependence on the number of agents n which grows only polynomially fast, leading to reasonable scaling for our protocol. The sieve constant of a graph, a new measure of connectivity we introduced, played a central role in our analysis.

Our research points to a number of intriguing open questions. Our results are for undirected graphs and it is unclear whether there is a learning protocol which will achieve similar bounds (i.e., a learning speed which depends only polynomially on n) on directed graphs. It appears that our bounds on the sieve constant given on Theorem 4 are loose by an order of magnitude (when compared with examples in Section 3) so that the learning speeds we have presented in this paper could potentially be further improved. In particular, comparing the scalings derived in Section 3 with the worst-case bound on the sieve constant given by Theorem 1.3 raises the possibility that Theorem 1.3 may not be tight. Moreover, it is further possible that a different protocol provides a faster learning speed compared to the one we have provided here.

Finally, and most importantly, it is of interest to develop a general theory of decentralized learning capable of handling situations in which complex concepts need to be learned by distributed network subject to time-varying connectivity and intermittent arrival of new information. Consider, for example, a group of UAVs all of which need to learn a new strategy to deal with an unforeseen situation, for example, how to perform formation maintenance in the face of a particular pattern of turbulence. Given that selected nodes can try different strategies, and given that nodes can observe the actions and the performance of neighboring nodes, is it possible for the entire network of nodes to collectively learn the best possible strategy? A the-

ory of general-purpose decentralized learning, designed to parallel the theory of PAC (Provably Approximately Correct) learning in the centralized case, is warranted.

REFERENCES

[1] D. Acemoglu, M. Dahleh, I. Lobel, A. Ozdaglar, "Bayesian learning in social networks," *Review of Econoics Studies*, vol. 78, no. 4, pp. 1201-1236, 2011.

[2] S. Adlakha, R. Johari, G. Weintraub, and A. Goldsmith, "Oblivious equilibrium: an approximation to large population dynamic games with concave utility," *Game Theory for Networks*, pp. 68-69, 2009.

[3] R. Alur, A. D'Innocenzo, K. H. Johansson, G. J. Pappas, G. Weiss, "Compositional modeling and analysis of multi-hop control networks," *IEEE Transactions on Automatic Control*, vol. 56, no. 10, pp. 2345-2357, 2011.

[4] M.-F. Balcan, F. Constantin, G. Piliouras, and J.S. Shamma, "Game couplings: Learning dynamics and applications," *Proceedings of IEEE Conference on Decision and Control*, 2011.

[5] B. Banerjee, J. Peng, "Efficient no-regret multiagent learning," *Proceedings of The 20th National Conference on Artificial Intelligence,* 2005.

[6] S. Boyd, P. Diaconis, and L. Xiao, "Fastest mixing markov chain on a graph," *SIAM Review*, 46(4):667-689, 2004.

[7] M. Bowling, "Convergence and no-regret in multiagent learning," *Advances in Neural Information Processing Systems*, pp. 209, 2005.

[8] G. Brightwell, P. Winkler, "Maximum hitting time for random walks on graphs," *Random Structures and Algorithms*, vol. 1, no. 3, pp. 263-276, 1990.

[9] G.C. Chasparis, J.S. Shamma, and A. Arapostathis, "Aspiration learning in coordination games", *Proceedings of the 49th IEEE Conference on Decision and Control*, 2010.

[10] G.C. Chasparis, J.S. Shamma, and A. Rantzer, "Perturbed learning automata for potential games," *Proceedings of the 50th IEEE Conference on Decision and Control*, 2011.

[11] K.L. Chung, "On a stochastic approximation method," *Annals of Mathematical Statistics,* vol. 25, no. 3, pp. 463-483, 1954.

[12] C. Claus, C. Boutellier, "The Dynamics of reinforcement learning in cooperative multiagent systems," *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, 1998.

[13] I. D. Couzin, J. Krause, N. R. Franks and S. A. Levin, "Effective leadership and decision-makingin animal groups on the move," Nature, 433, 2005.

[14] R. Gopalakrishnan, J. R. Marden, and A. Wierman, "An architectural view of game theoretic control," *ACM Sigmetrics Performance Evaluation Review*, 38(3), 2011.

[15] V. Guttal and I. D. Couzin, "Social interactions, information use, and the evolution of collective migration," *Proceedings of the National Academy of Sciences*, 107(37), 2010.

[16] H. J. Landau, A. M. Odlyzko, "Bounds for eigenvalues of certain stochastic matrices," *Linear Algebra and Its Applications,* vol. 38, pp. 5-15, 1981.

[17] N. E. Leonard, T. Shen, B. Nabet, L. Scardovi, I. D. Couzin and S. A. Levin, "Decision versus compromise for animal groupsin motion", *Proceedings of the National Academy of Sciences*, 109(1), 2012.

[18] M. L. Littman, "Markov games as a framework for multi-agent reinforcement learning, *Proceedings of the Eleventh International Conference on Machine Learning*, 1994.

[19] S. Martinez, F. Bullo, J. Corts, E. Frazzoli, "On synchronous robotic networks - Part I: Models, tasks and complexity," *IEEE Transactions on Automatic Control*, vol. 52, no. 12, pp. 2199-2213, 2007.

[20] S. Mannor and J. Shamma, "Multi-agent learning for engineers," *Artificial Intelligence*, vol. 171, no. 7, pp. 417-422, 2007.

[21] J.R. Marden, G. Arslan, J.S. Shamma, "Cooperative control and potential games," *IEEE Transacions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 39(6), pp. 1393-1407, 2009.

[22] A. Nedic, A. Olshevsky, A. Ozdaglar, J.N. Tsitsiklis, "On Distributed Averaging Algorithms and Quantization Effects," *IEEE Transactions on Automatic Control,* vol. 54, no. 11, pp. 2506-2517, 2009.

[23] A. Nedic and A. Ozdaglar, "Distributed Subgradient Methods for Multi-agent Optimization," *IEEE Transactions on Automatic Control*, vol. 54, no. 1, pp. 48-61, 2009.

[24] L. Panais, S. Luke, "Cooperative multi-agent learning: the state of the art," *Autonomous Agents and Multi-Agent Systems*, vol. 11, no. 3, pp. 387-434, 2005.

[25] B. Polyak, *Introduction to Optimization*, New York: Optimization Software Inc, 1987.

[26] T.W. Sandholm, R.H. Crites, "Multiagent reinforcement learning in the iterated prisoner's dilemma," *Biosystems*, vol. 37, no. 1-2, pp. 147-166, 1996.

[27] P.Stone, M. Veloso, "Multiagent systems: a survey from a machine learning perspective," *Autonomous Robots*, vol. 8, no. 3, pp. 345-383, 2000.

[28] K. Srivastava and A. Nedic, "Distributed asynchronous constrained stochastic optimization," *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 4, pp. 772-790, 2011.

[29] M. Tan, "Multi-agent reinforcement learning: independent vs. cooperative agents," *Proceedings of the Tenth International Conference on Machine Learning*, 1993.

[30] C. Torney, Z. Neufeld and I. D. Couzin, "Context-dependent interaction leads to emergentsearch behavior in social aggregates," *Proceedings of the National Academy of Sciences*, 106(52), 2009.

[31] B. Touri and A. Nedic, "On existence of a quadratic comparison function for random weighted averaging dynamics and its implications," *Proceedings of the 50th IEEE Conference on Decision and Control,* 2011.

[32] F. Yousefian, A. Nedic, U.V. Shanbhag, "On stochastic gradient and subgradient methods with adaptive steplength sequences," *Automatica*, vol. 48, no. 1, pp. 56-67, 2012.